# Lessons from Implementation of Beacon Spacecraft Operations on Deep Space One

Robert Sherwood, Alan Schlutsmeyer, Miles Sue, E. Jay Wyatt

Jet Propulsion Laboratory
California Institute of Technology
4800 Oak Grove Dr.
Pasadena, CA  91109
818-393-5378
818-393-5244 FAX
rob.sherwood@jpl.nasa.gov

*Abstract* - A new approach to mission operations has been flight validated on NASA's Deep Space One (DS1) mission that launched in October 1998.  The beacon monitor operations technology is aimed at decreasing the total volume of downlinked engineering telemetry by reducing the frequency of downlink and the volume of data received per pass.  Cost savings are achieved by reducing the amount of routine telemetry processing and analysis performed by ground staff.  With beacon monitoring, the spacecraft will assess its own health and will transmit one of four sub-carrier frequency tones to inform the ground how urgent it is to track the spacecraft for telemetry.  If all conditions are nominal, the tone provides periodic assurance to ground personnel that the mission is proceeding as planned without having to receive and analyze downlinked telemetry.  If there is a problem, the tone will indicate that tracking is required and the resulting telemetry will contain a concise summary of what has occurred since the last telemetry pass.  The Beacon technology has been proven successful on DS1 through a series of tone tests and data summarization experiments.  This collection of experiments was called the DS1 Beacon Monitor Experiment or BMOX.  Still there are important lessons learned from this experiment that can be applied to future spacecraft missions.

## TABLE OF CONTENTS

## 1. INTRODUCTION

The budget environment that has evolved since the advent of NASA's Faster, Better, Cheaper initiative has caused mission risk policies and mission designs to change in ways that have been conducive to the inception of new operations concepts and supporting technologies.  Such was the case when the beacon monitor concept was conceived to enable a mission to Pluto to be achieved within the budget constraints passed down from NASA.  The technology was accepted into the New Millennium Program and baselined for flight validation on the Deep Space One Mission.  As the technology was being developed for DS1, the NASA community has expressed a growing interest and acceptance of adaptive operations and onboard autonomy.

In traditional mission operations, the spacecraft receives commands from the ground and in turn transmits telemetry in the form of science or engineering data.  With beacon monitoring, the spacecraft sends a command to the ground that instructs the ground personnel how urgent it is to track the spacecraft for telemetry.  There are only four such commands.  Thinking of beacon operations in this way creates a paradigm shift over the way we traditionally approach operations.  Also, it is very important to not think of the tone message as just a little bit of telemetry.  If one does this, it is easy to make the argument that a little more telemetry is better.  Our approach is one where telemetry is only transmitted when it is necessary for ground personnel to assist the spacecraft or otherwise very infrequently if the spacecraft is fortunate enough to go long periods (a month or so) without requiring ground assistance.  When telemetry tracking is necessary the intelligent data summaries contain the most relevant information to provide full insights into spacecraft activities since the last contact.  The key challenge has been to develop an architecture that enables the spacecraft to adaptively create summary information to make best use of the available bandwidth as the mission progresses such that all pertinent data is received in one four to eight hour telemetry pass.

The primary components of the technology are a tone messaging system, AI-based software for onboard engineering data summarization, a ground visualization system for telemetry summaries, and a ground response system. Beacon tone operations can be used to lower the cost of operating space missions while simultaneously decreasing their risk. The concept involves a paradigm shift from routine telemetry downlink and ground analysis to on-board health determination and autonomous data summarization. Ion Propulsion missions gain an added advantage of power savings from reduced telemetry downlinks and the associated increased thrusting time. Beacon operations will enable more of the smaller, more frequent missions that NASA is planning for the early part of the next millennium. This paper will document the results of the beacon experiment on DS1. In addition, the paper will include a description of the Beacon monitor concept, the trade-offs with adapting that concept as a technology experiment, and our lessons learned during the mission. Applicability to future missions will also be included.

## 2. DS1 BMOX SUBSYSTEMS

It was required that two subsystems be designed and developed to implement the desired functionality for the DS1 experiment. These are, in fact, standalone innovations. Although they are being presented here primarily in support of cruise phase operations, there has also been interest in applying these technology components to other domains. Other potential applications include using in-situ beacons at Mars, adapting tone messaging and summarization to earth orbiters, using beacons for science event detection and notification, and in utilizing the tone system to reduce mission risk due to spacecraft operability constraints.

*Tone system*

There are four tone signals and each uniquely represents one of the four urgency-based beacon messages. The DS1 tone definitions are summarized in Table 1. These tones are generated as the spacecraft software reacts to real-time events.

Urgent Beacon tones on DS1 are sent when the spacecraft fault protection puts the spacecraft in standby mode. This condition occurs when the fault protection encounters a fault that it cannot correct. Standby mode halts the current command sequence, including IPS thrusting. During the DS1 tone experiment, the Beacon tone can be sent regularly at a prescheduled time, i.e., 30 to 60 minutes per day. The tone cannot be operated continuously because DS1 requires as much power as possible for IPS thrusting and the tone transmission uses some of the thrusting power. Routine operational use of the beacon monitor system is currently being explored for the DS1 extended mission, scheduled to begin in September of 1999.

Table 1. Tone Definitions

| Tone | Definition |
|---|---|
| Nominal | Spacecraft is nominal, all functions are performing as expected. No need to downlink engineering telemetry. |
| Interesting | An interesting and non-urgent event has occurred on the spacecraft. Establish communication with the ground when convenient. Examples: device reset to clear error caused by single event upset due to cosmic particle, other transient events. |
| Important | Communication with the ground needs to be achieved within a certain time or the spacecraft state could deteriorate and/or critical data could be lost. Examples: memory near full, non-critical hardware failure. |
| Urgent | Spacecraft emergency. A critical component of the spacecraft has failed. The spacecraft cannot autonomously recover and ground intervention is required immediately. Examples: Propulsion or power system electronics failure |
| No Tone | Beacon mode is not operating, spacecraft telecom is not Earth-pointed or spacecraft anomaly prohibited tone from being sent. |

It is important to communicate the urgency of ground response using a telecommunications method that has a low-detection threshold and short detection times. Ease of detection translates to lower cost operations. The signal structure is shown in Figure 1. Each message is represented by a pair of tones centered about the carrier frequency. Tones are generated by phase-modulating the RF carrier by a squarewave subcarrier using a 90 degrees modulation angle. The carrier ($f_c$) is completely suppressed. The resulting downlink spectrum consists of tones at odd multiples of the subcarrier frequency above and below the carrier. Four pairs of tones are needed to represent the four possible messages.
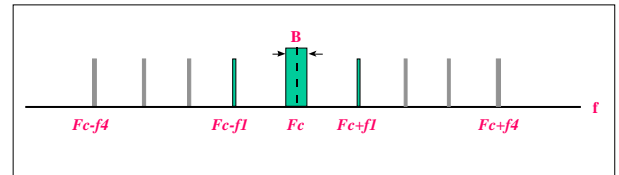


Figure 1. Tone Signal Structure
B=Frequency uncertainty    Fc=Carrier frequency
fi=Subcarrier frequency for the i[th] message

The goal is to reliably detect the monitoring messages with zero dB-Hz total received signal-to-noise-spectral-density ratio (Pt/No) using a 1000 second observation time. Future missions are assumed to carry a low-cost auxiliary oscillator as a frequency source, instead of a more expensive, ultra-stable oscillator. The downlink frequency derived from the auxiliary oscillator is not precisely known due to frequency

drifts caused by on-board temperature variations, aging, and uncorrected residual Doppler frequency. In addition, the downlink frequency also exhibits short-term drift and phase noise. These factors were taken into consideration in the design of the monitoring signal detector.

*Onboard summarization system*

If the beacon tone indicates that tracking is required, the onboard summarization system provides concise summaries of all pertinent spacecraft data since the previous contact. The summarization system performs three functions: data collection and processing, mission activity determination, and episode identification. The data collection subroutine receives data from the engineering telemetry system via a function call and applies summary techniques to this data, producing summary measures for downlink to the ground. The mission activity subroutine determines the overall spacecraft mode of operation. This determination is used to choose the appropriate data and limits monitored by the episode subroutine. The mission activity is intended to be exclusive. When a new mission activity starts, the previous mission activity is assumed to have ended. The episode subroutine combines summary and engineering data received internally from the data collection subroutine with the mission activity received from the activity subroutine and compares the data with mission activity specific alarm limits. It is necessary to use the mission activities to determine which data to use for episode identification and to identify the limits of these data. If the limit is exceeded, the subroutine spawns a new episode and collects past relevant data from the data collection subroutine. The past data collected will be one-minute summaries that go back in time as far as the user has defined. (So a five-minute episode would contain summaries starting five minutes before the episode to five minutes after the episode.) At the end of the episode, the subroutine outputs data to the telemetry subsystem for downlink.

Table 2. Summarization Telemetry Packets

| Telemetry Name | Description | Output Frequency |
|---|---|---|
| Activity | Current value of mission activity | Output on change |
| Data Sample | Records a snapshot of every raw and summarized data channel | Regular interval, i.e., 15 min. |
| Episode Summary | Records general data about an out-of-limits data condition called an "episode" | One per episode |
| Episode Channel | Records specific data about a single data channel's behavior during an episode | One or more per episode |
| Tone Change | Current state of the beacon tone | Output on tone change |
| Channel Summary | Summary data about a single data channel's behavior since the last downlink | One for each channel out of limits |
| User Summary | A user-specified packet containing raw and/or summarized data | Duration user-specified |

Three different types of summarized data are produced onboard: overall performance summary, user-defined performance summary, and anomaly summary. Six different telemetry packets have been defined to contain this information. (See Table 2.) Taken as a whole, the telemetry packets produce summary downlinks that are used to enable fast determination of spacecraft state by ground personnel. The performance summaries are generated at regular intervals and stored in memory until the next telemetry ground contact. They are computed by applying standard functions, such as minimum, maximum, mean, first derivative, and second derivative, to the data. User-defined summary data can provide detailed information on a particular subsystem and are created at the user's discretion. Anomaly summary data (episodes) are created when the raw and summarized data violate high or low limits. These limits are determined by the subsystem specialist and stored in a table on-board the spacecraft. The limit tables are based on the current mission activity.

The software also has the capability to use AI-based envelope functions instead of traditional alarm limits. This system, called ELMER (Envelope Learning and Monitoring using Error Relaxation), provides a new form of event detection will be evaluated in addition to using the project-specified traditional alarm limits. Envelope functions are essentially adaptive alarm limits learned by training a neural network with nominal engineering data. The neural net can be onboard or on the ground. For DS1, envelope functions are trained on the ground and then uploaded to the spacecraft. DS1 spacecraft fault protection will only be based on project-specified static alarm limits but the summary data can be generated based on the adaptive limits.

## 3. TONE SYSTEM RESULTS

A series of experiments were run to test the end-to-end tone delivery system. These experiments were designed to incrementally test additional capability for the Beacon tone system. Prior to launch, the ability of the Small Deep Space Transponder (SDST) to generate Beacon tones was tested by the telecom engineers. A similar test was performed on the spacecraft several times after launch. This test was called "X-tone" because it tested the capability to send the Beacon tones using X-band transmission. The X-tone test, expanded to use a series of tones to test the ground detection system, was repeated several times throughout March and April 1999.

The ability of the software to select tones and transmit them in DS1 telemetry was tested on February 26, 1999. This test, called b-tone, consisted of ground commands that set the Beacon tone during a downlink pass. The tone was verified in regular DS1 telemetry but was not transmitted to the tone detector. Each tone was verified during the b-tone test. In addition, the tone-reset command was tested.

The next test to run on-board DS1 was the b-transmit test. This test involved setting the Beacon tone using information from the software on board, then transmitting the tone using

the SDST. The tone was received at the DSS-13 antenna and forwarded to the tone detector at JPL. No advance knowledge of the commanded tone was given to the ground detection engineer. After the tone was detected, it was delivered to other members of the Beacon Team in an email message. The b-transmit test was run three times in April 1999.

The last tone test to be run was the Ka-tone test. This test was identical to the X-tone test except that it used the Ka-band transmitter to send the Beacon tone. This test was run in April 1999.

## 4. DATA SUMMARIZATION RESULTS

The data summarization component of BMOX was first turned on February 19, 1999. The Beacon Team determined the limits applied to the engineering data for testing the summarization capability. The limits were set just outside of the minimum and maximum value seen for the data since launch. Shortly after the first turn-on several of the data channels went into episode (out-of-limits) condition. Upon further inspection, it was determined that many limits were based on engineering units (EU), but much of the data was being stored using data numbers (DN) in the on-board engineering and housekeeping telemetry system (EH&A). The data summarization was turned off after several hours and the initialization file (also called sampler init file or SIF) was updated with DN based limits.

The data summarization was turned back on March 8th for several hours. A few channels went into alarm, but the number was reduced from the previous test. Inspection of the data revealed negative values for some eight-bit sensors. This was impossible because all eight-bit sensors should range from 0 to 255. After careful debugging in the DS1 test bed, an error was found in the DS1 flight software. It was discovered that when data are passed from the originator to EH&A, EH&A converts the data to its own internal double precision format as though it were 8 bits and signed. This results in the values from 0 to 127 being represented correctly, and the values from 128 to 255 being represented as -128 to -1, respectively. EH&A apparently does not have a data type code for unsigned 8-bit integers. The effect of this problem was that limits were harder (and sometimes impossible) to specify. With a new set of rules, we were able to create a SIF that would work around this problem for some of the data. If both high and low limits were 128 or greater, they had to be converted by subtracting 256. However, if the low limit was 127 or less and the high limit is 128 or greater, the limits won't work. Sensor values with both limits less than 127 could remain unchanged. With these rules, we created another SIF and uploaded it to DS1. Data summarization was restarted on March 22, 1999. Everything appeared to operate correctly in data summarization. A few data channels went into episode condition. It was determined that temperature sensors were

drifting colder due to DS1 moving away from the sun. The limits were updated and a new SIF was uplinked.

Data summarization ran smoothly on and off during the month of April and May with minor modifications to the SIF due to noisy channels. During this period, a new version of the Beacon FSW was developed and tested. This version included a work-around for the limitation of EH&A data described above. In addition, other new features were added:

- ♦ The criteria for determining mission activity was parameterized in the SIF
- ♦ Episodes will now end if a new SIF is loaded
- ♦ Additional protection for divide-by-zero conditions
- ♦ SIF's can now be loaded from EEPROM or RAM
- ♦ User data packets can now have start and stop times associated with them

The new version was started up on June 15th, 1999. A new SIF was included with limits determined by the DS1 spacecraft engineers. Since that time, data summarization has needed a few updates due to false alarms. There are several reasons for these false alarms. The Beacon flight software (FSW) is able to sample the data once per second. This is a much higher rate than the data sent to the ground for analysis. Because of the higher rate, the FSW is able to see events that are normally missed on the ground. We have confirmed these events by correlating with the fault protection monitors that capture maximum excursions on some engineering data.

Table 3. Summary of Engineering Data Monitored

| Subsystem | Number of Channels |
|---|---|
| Attitude Control | 8 |
| Fault Protection | 1 |
| Navigation | 1 |
| Other | 2 |
| Power | 22 |
| Propulsion | 1 |
| Telecommunications | 6 |
| Temperature (all subsystems) | 35 |

Another reason for false alarms has been activities such as optical navigation (OPNAV's) that move power and thermal sensors outside their normal ranges. The subsystem engineers respond, "yes, these events take the sensors outside their normal ranges, and yes, this is expected behavior." So where does the Beacon Team set the limits? Since the Beacon data summarization is context sensitive, we could create a new "mission activity" for OPNAV's with it's own set of limits. An OPNAV activity consists of several spacecraft turns, with picture taking occurring at each target. This is similar to a maneuver. With this mind, we have changed our mission activity determination criteria for maneuvers to include optical navigation activities. This will also make our maneuver activity determination more

robust. Prior to this change, we were only changing to maneuver activity when DS1 was actually firing thrusters to change the velocity. Maneuvers involve turning to a thrusting attitude, and turning back after the thrusting. Now, the maneuver activity includes these turns and their respective settling times as well. This makes sense because it is during this entire period that power and thermal sensors may deviate from their nominal cruise values. This change was uplinked in early September 1999. A summary of this list is contained in Table 3.

Beacon data summarization has been an evolving process requiring several limit refinements from the spacecraft team. This should be expected in the development of any data summarization system. This process is very similar to when any new mission launches. For the first several months, ground alarms are updated as the flight team learns about how the spacecraft really operates. The ground testing activities give a good first cut at setting alarm levels, but the spacecraft never operates exactly as it did in test. Implementing context sensitive limits is a similar process. We are no longer setting the engineering data limits based on the worse case. Now we can look at the worse case based on the spacecraft activities. This should ensure more accurate discovery of anomalies.
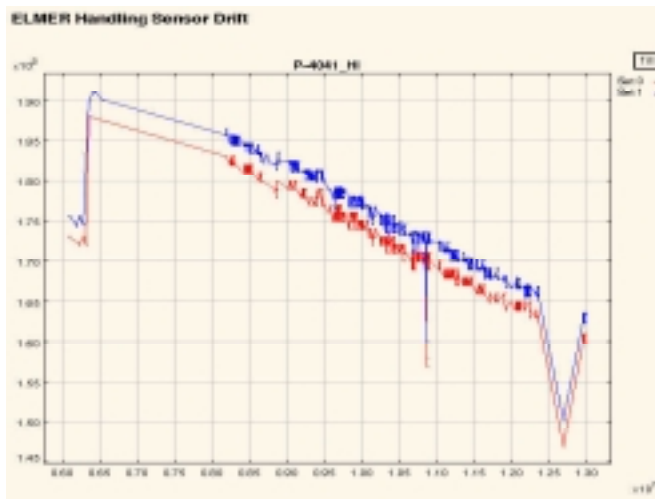


Figure 2. Tracking of adaptive alarm limit to DS1 solar array temperature

One activity that produced important results involves analyzing summary system performance on DS1 anomalies to date. Although limited in its capabilities due to on-board memory restrictions, preliminary results when running ELMER on historical data are showing that adaptive alarm thresholds can track gradual trending of sensor data much tighter than the current DS1 static alarm limits. We see this in monitoring the gradual drift in eight solar array temperatures sensors, one of which is shown in Figure 2. In comparing traditional limits with ELMER limits during the 81 days of operations, we see that ELMER limits track

actual spacecraft performance much more precisely than static limits, which would be off the scale of this chart.

Another validation exercise has confirmed that summarization can capture subtle, yet important spacecraft episodes. In ground tests, ELMER detected an unexpected heater turn-on that occurred when the solar panels went off-axis during a spacecraft maneuver. Since ELMER trains across multiple parameters using nominal data, the summarization system detected this event without explicit a priori knowledge of the scenario. This data is shown in Figure 3.
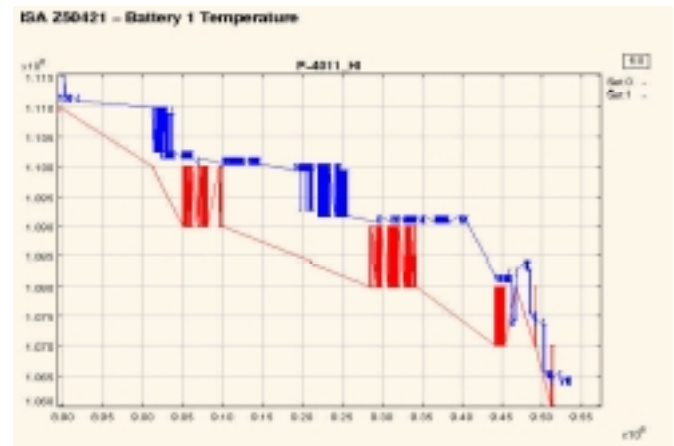


Figure 3. Battery Temperature Episode Detection

ELMER has been running on-board with only 10 sensors, all temperature. This limitation is primarily due to limited on-board memory. There have only been three ELMER limit violations (episodes) during the primary mission. Two have occurred during OPNAV events and can be explained by the temperature excursions associated with spacecraft turns. These are basically "false alarms." The third episode has not yet been explained. The ELMER limit functions were developed after training on data from the first four months of the mission. We hope that additional training on spacecraft data since February will correct these false alarms in extended mission. There will be additional ELMER limit functions added in extended mission as well.

## 5. LESSONS LEARNED

*Ion Propulsion Missions*

The utilization of the ion propulsion system (also called solar-electric propulsion) on DS1 offers an additional advantage in using Beacon monitoring. The IPS provides continuous thrust for much of the cruise phase. The operational margin for IPS thrusting represents the duration for which IPS could be off and still allow the spacecraft to reach the target asteroid. Due to the low thrust associated with IPS and because actual thrusting did not start until several weeks after launch, the operational margin is only a few weeks. Telemetry downlink passes are becoming less frequent as the DS1 mission progresses. Eventually, there

will only be one telemetry pass per week. If the spacecraft experiences a problem that requires the standby mode, the IPS engine will be shut down. It could be up to one week before the flight team has visibility to that standby mode. Using the Beacon tone system during the periods between scheduled telemetry downlinks can be a cost-effective way to decrease mission risk because it reduces the likelihood of losing thrusting time and not making the intended target. Other future IPS missions have taken note of this fact and requested Beacon tone services to lower their mission risk.

*Software Testing*

It was decided to redesign the DS1 flight software about 18 months before launch. This decision greatly compacted an already full schedule to complete the software. As a result, the testing of all non-essential software functions was delayed until after launch. The Beacon experiment was considered a non-essential piece of software and therefore was only tested pre-launch for non-interference with the other flight software. In post launch testing, a few problems were discovered that prevent us from starting the Beacon software until a new version could be uploaded. These problems related to differences between the flight hardware-based testbed and a simulated hardware testbed. This is the age-old lesson learned of performing system testing on the software prior to use. But even beyond that, it is important to run tests on the actual hardware-based testbed. Unfortunately, the DS1 schedule would not allow us to do this until post launch.

*Fault Protection Integration*

Before the software redesign, the Beacon software was tightly integrated with the DS1 fault protection software. The decision was made after the redesign to de-couple the two pieces of software. Previously, the fault protection monitors triggered the Beacon tones. After the redesign, the mapping of faults to tones was performed using two different methods. All spacecraft standby modes are now mapped to the urgent Beacon tone. The interesting and important Beacon tones are mapped using Beacon software determined limits. Decoupling the fault protection software from the Beacon software gives us maximum flexibility to determine what sensors to monitor. Unfortunately, our algorithms for determining faults are not nearly as sophisticated as the fault protection monitors. These monitors can look at many different values based on conditional logic before determining what fault has occurred. Future spacecraft designed to use Beacon operations should plan on completely integrating the Beacon tone software with the fault protection software.

*Beacon Signal Frequency Stability*

The signals used for Beacon monitor are characterized by three things: (1) the signal strength can be extremely low, (2) the initial tone frequencies, which are derived from an on-board auxiliary oscillator, are not known exactly, and (3) the tone frequencies are constantly drifting. The tone detector is designed to detect these types of signals with a high-level of confidence. The maximum frequency uncertainty and the maximum frequency drift rate for the tone detector were established using a Galileo spare transponder. An operational issue was encountered with the DS1 Beacon experiment: how and to what extent can we stabilize the temperature of the auxiliary oscillator before the start of a Beacon pass? Stabilizing the temperature will reduce the frequency uncertainty and frequency drift, making it easier for the tone detector to detect the Beacon signal. Based on data provided by the DS1 telecom personnel, the auxiliary oscillator temperature can undergo a wide range of changes after an OPNAV maneuver. This results in a very large frequency uncertainty and a very high rate of change ($>6$ Hz/sec), both of which would exceed the limits of the tone detector (when the signal level is low).

One solution to overcome the OPNAV-related problem is to wait for the transponder temperature to stabilize. Studies by the DS1 telecom personnel indicated that about four hours are needed for the transponder temperature to stabilize after running the OPNAV activity. This operational constraint would not have much impact on the spacecraft and is believed to be the simplest, lowest-cost solution to this problem. We recommend this procedure to improve weak-signal detection for DS1 and future missions using Beacon Monitor.

During the DS1 tone experiments, the initial frequency uncertainty was much larger than expected. A bias was manually introduced to keep the received signal in the recorded band. Without the bias, the frequency might be outside the recorded band. In an automated detection mode, it is necessary to record at least 3 times the current bandwidth, unless a better way to predict the frequency can be found. One possibility is to make use of the Auxiliary Oscillator Frequency vs. Temperature calibration table to improve frequency prediction.

*Downlink Carrier Phase Noise*

Post analyses of the received signal frequency indicated that the phase noise of the downlink carrier was fairly significant. This would result in detection loss. Analyses should be performed to estimate the impact of this phase noise on detector performance and factor this into future detection experiments.

*Spacecraft Clock Accuracy*

During one of the experiments, it was observed that the actual tone switching times did not seem to agree exactly with the predicted switching times. This led to the discovery by the DS1 team that there was an error of 18-19 seconds in the on-board spacecraft time to earth time conversion.

*DSN Equipment Issues*

A couple of tone passes were not successful due to the DSN station's (DSS-13) weather and equipment. In one

experiment, the spacecraft started transmitting tones before it rose above the horizon of DSS13. In another case, a scheduled pass was cancelled due to spacecraft activities. While the overall tone experiments have been very successful, future experiment plan should allow for this kind of contingency.

### Beacon Operations Paradigm

The Beacon software makes determinations of spacecraft anomalies. The data summarization component of Beacon attempts to summarize related data from these anomalies. These determinations are based upon high and low limits on sensor data. It is important to involve the spacecraft subsystem engineers in the determination of which data to monitor and the setting of the limits on these data. They are the personnel most familiar with the operational characteristics of each subsystem and therefore should be determining interesting and fault conditions for their subsystem. Also, by involving them in the data summarization definition, they will become better acquainted with the Beacon software and will be more inclined to use it during crisis situations.

Ground alarm limits on telemetry are generally set using the worse possible state of each data channel. This practice can hide problems with the spacecraft if the alarm limits are set at wide boundaries. Beacon data summarization offers context sensitive limits. In the case of DS1, limits can be set for cruise, downlink, IPS thrusting, maneuver, and standby modes. Spacecraft operations personnel are not used to working with summarized engineering telemetry or context sensitive limits. When asked for data limits, we generally received one set of limits and were told to apply them to all mission activities. Setting limits like this does not utilize the capabilities of the Beacon data summarization. For future implementations of Beacon, it will be important to educate the flight team about Beacon's capabilities early in mission design. Beacon data summarization should also be used during spacecraft testing to familiarize operators with the technology. This will help ensure reliance on Beacon data during the mission.

### Systems Engineering

As previously mentioned, there were problems with false episode alarms due to mission activities such as Optical Navigations, camera calibrations, etc. It is important to carefully define each of the mission activities and how they are related to engineering data. In the DS1 case, we had defined the maneuver activity to only occur when the thrusters were firing. Since maneuvers also involved turning the spacecraft, it was important to include all events that turned the spacecraft in our maneuver mission activity criteria. Once mission activities are carefully defined, then episode limits for those activities can be developed.

### Other Possible Implementations

Earlier it was stated that the lack of a Beacon tone implied there was a problem with the telecommunication system or Beacon software. It's also possible to consider non-detection a good response since an autonomous spacecraft may be doing something more important than just telling the ground it's OK, but that is not true indefinitely. If you don't detect the spacecraft for some number of days then you have a problem. In other words, time since previous tone and tone history are both necessary to interpret the Beacon tone.

There is another proposed Beacon concept for an earth trailing spacecraft (SIRTF) that involves using one tone. SIRTF plans to track every 12 hours, but would like to have Beacon tracking every 2 hours. The idea is that the spacecraft would only send a Beacon tone if it had a problem. The possible Beacon detections are 1) help tone, or 2) no detection. Normally the spacecraft would be busy doing observations, but if it had a problem it would turn to earth point and start transmitting a carrier signal. This Beacon signal could shorten the anomaly response time from 12 hours to a maximum of 2 hours. This requires no modification to the already designed spacecraft since there is no need to distinguish fine levels of urgency. SIRTF management considers this important because their design does not include a transponder that supports Beacon tones. There is one drawback with this operation. When the tone detector fails to detect a Beacon signal, one can not tell whether (1) the spacecraft is fine and no Beacon has been transmitted, or (2) the spacecraft has an anomaly and fails to transmit.

## 6. OPERATIONAL EFFECTIVENESS

The Beacon Operation Monitor Experiment really was just an experiment to test the pieces of a new technology. DS1 never relied on the Beacon tone or data summarization for the operations of the spacecraft. Beacon was not given many spacecraft resources or time because there were 11 other experiments to test. BMOX never really was able to get to the point where a true end-to-end long-term test of the technology could be performed. Despite this fact, we were able to discover some additional innovations that would make full operational use of Beacon easier.

### Data Processing Issues

Beacon summary data was delivered to the Beacon Team through an automated batch script that queried the data each night. The data was placed in a public directory and then processed by the Beacon Team the next morning. The processing was a simple task, but was not automated because data summarization was frequently turned off for days to weeks at a time. During DS1 extended mission, data summarization should be on continuously and therefore the data processing should to be automated.

The database used to store Beacon summary data was created specifically for the Beacon Task. Because summary data is not easily formatted for commercial databases, we decided to develop our own database. In hindsight, we believe this was the wrong decision. It has been very difficult to maintain a custom database. The users do not have good visibility into the database if the tools are not working correctly. Changes to the database take a programmer to change the code instead of running a tool that would be provided with a commercial database. In addition, commercial databases have built in query features that are easy to set-up and use. There were instances in which we were asked for data and were not able to provide it in a timely fashion. We also were not able to do custom queries such as a query for all episodes involving a specific channel. The limitations of using a custom database hindered the operational effectiveness of Beacon.

*Data Summarization Software Enhancements*

The data summarization software was not relied upon for determining spacecraft state. Although the algorithms and returned summary data seemed adequate, there were several suggestions made by the Beacon personnel and flight team for further enhancements. Some of these suggestions will be incorporated into the next (M7) version of the flight software to be uploaded during DS1 extended mission operations.

The episode data was lacking depth because it only provided ten samples; each separated by two minutes. The long time between samples was set to ensure that Beacon summary data would not overflow the telemetry buffer in the event of repeated episodes on a single channel. For the M7 version of the software, we are changing the number of samples to 20, and allowing the user to set the number of times a channel can go into episode before it stops producing episode packets. With these changes, we can set the sample interval much shorter. In fact, we plan on using a six-second-sample interval. This will give us much more visibility on the episodes while not overloading the telemetry buffer with false alarms. We considered making a change to add all data on change to episodes, but the DS1 project only wanted very minor software changes in M7.

During the course of operations, the initialization file with the episode limits was changed and uplinked many times. Many times the changes only involved one or two limits in the file. Because the file is on the order of 15 kilobytes, there were periods of low communications bandwidth when it would take several minutes to uplink the file using the low gain antenna. Operationally, it would have been much easier if we had a capability to update limits without sending the entire initialization file.

The flight team made a few suggestions for improving the usefulness of the summary data. We have already implemented derivative summary functions, but one of the subsystems suggested that integrals be added to the summary functions. Several other flight team members suggested adding different persistence for each episode limit check. Currently, we have a global persistence parameter that applies to all episodes. This change will be implemented in our M7 software release. Another suggestion was to add a sample rate to user performance packets.

Two capabilities that fault protection monitors have that should be present in Beacon are conditional monitors and maximum excursion tracking. Conditional monitors enable the user to check multiple sensors based on the values of the sensors. The DS1 fault protection software also has the capability to track and save the minimum and maximum values for sensors. The summarization software will only track these values if the sensor goes into an episode condition. This may be important data for future missions relying on summary data even though the sensors are not outside their limits. As mentioned in the Lessons Learned section, there should be tighter integration between the Beacon software and the fault protection software.

*Reporting Results to the Flight Team*

We developed a set of tools for examining the summary data. These tools were only located on the Beacon Team workstation. Since launch we have developed some web-based tools to access the summary data. These tools have made it easier to report the results to the flight team, but are very limited in their capabilities. We are going to improve these tools during extended mission. Our goal is to make the data easily accessible to the flight team users. Easy access to the Beacon data is very important for making the technology operationally effective but unfortunately was not available during the DS1 primary mission.

*Cost Savings from Using Beacon*

Part of our future work in Beacon technology involves infusing the Beacon technology into DS1 mission operations as an end-to-end system. Technology infusion is not an easy task and traditionally has not been done well. DS1 will benefit from this work by reducing the amount of tracking time used.

In extended mission, DS1 will have two tracking passes per week, an 8-hour high gain pass on Monday's and a 4-hour mid-week pass to check spacecraft status. Utilizing Beacon, the DS1 project will not have to use a 4-hour mid-week carrier only DSN pass to check spacecraft status. They can use a 30 minute (or less) Beacon pass that actually gives them additional information over a carrier only pass. In addition, they can reduce the frequency of eight-hour telemetry passes and substitute 30-minute Beacon passes in their place. We have not yet determined how many 8-hour telemetry passes could be eliminated but DS1 expects it could be as many as every other pass. In this case, there would only be two eight-hour telemetry passes each month and four 30-minute Beacon passes each month. The overall

savings for this case are summarized in the Table 4. This results in savings of 30 hours of DSN tracking time or $18,248 per four-week period. This does not include the substantial savings of mission engineering labor costs of performing routine telemetry analysis.

Table 4. Tracking Cost Per Month
(34m BWG, 2 contacts per week, assuming reduction of two 8-hour telemetry passes per month)

| | DS1 Operations without Beacon | DS1 Operations with Beacon | Total Monthly Savings |
|---|---|---|---|
| 8-hour telemetry passes | $19,465 | $9,733 | |
| 4-hour carrier only passes | $9,733 | $0 | |
| Beacon tone passes | $0 | $1,217 | |
| **Total** | $29,198 | $10,950 | $18,248 |

The benefits of infusing a regular Beacon operation technology on DS1 are apparent in the cost savings of reduced DSN utilization. In addition, the four-hour mid-week passes are replaced with 30-minute Beacon passes that contain additional status information. Future missions will benefit from the experience of a flight mission using a regular Beacon tone for an extended period of time. This includes the experience of scheduling the DSN for Beacon operations as well as the success of the Beacon tone system in relaying the spacecraft status to the ground. New missions that could benefit from this technology include Pluto Express, Europa Orbiter, and MDS. Each of these missions is planning on using either part or all of the Beacon operations technology. The continuation of work on the Beacon technology by revising the operations concept will add value to these mission customers. In addition, we can fully develop the operations procedures for using the Beacon technology.

Demand-access scheduling of DSN antennas is another important feature of an operational Beacon system. Scheduling antennas based on demand rather than a pre-negotiated agreement is important to the success of this technology within the DSN. During the DS1 extended mission, we do not have the funding to demonstrate automated scheduling of antenna resources. If we receive a Beacon tone that requires contacting the DS1 spacecraft, we will have to manually request a station pass. Until the DSN changes their scheduling paradigm, it will be difficult to implement demand-access scheduling.

## 7. FUTURE WORK

The DS1 Beacon Monitor Experiment involved testing the functionality of an onboard tone system and data summarization capability. A total of twelve-tone experiments were conducted. For these experiments, a pre-selected tone (subcarrier) or a sequence of tones was uploaded to the spacecraft prior to the experiment. The tone was then detected on the ground and compared with the tone sent. Although this functionality was proven to be successful, the previous Beacon work was focused on validating the technology, not making it useful to the DS1 mission.

The end-to-end Beacon concept involves using the current spacecraft health to determine a tone and then transmit that tone to a ground station. The objective of this task is to adapt the end-to-end Beacon operations technology for infusion into the operations of a functional spacecraft. Although DS1 and Mission Data Systems (MDS) are eager to reduce operation costs using the Beacon operations technology, infusion into an operational mission still requires a technology push.

The DS1 extended mission will be used to test this technology. The current tracking plan for DS1 extended operations includes a once per week 8 hour DSN pass to send telemetry data to the ground. In addition, mid week 4-hour carrier only passes are planned to confirm operation of the ion propulsion system (IPS) thrusting. We propose using a Beacon tone pass in lieu of the mid-week pass. Using this mid-week Beacon pass may also reduce the need for weekly 8-hour telemetry passes.

End-to-end Beacon operations have never been performed on an operational spacecraft. Regular operations using a Beacon tone will have implications on the scheduling of DSN tracking for DS1. Solving these scheduling problems will be an added benefit of this task and will help enable future missions to use Beacon operations.

It will be important to automate tone detection operations to support DS1 extended mission. This will involve implementing algorithms for unattended operations to support operational use of this technology by DS1. Since we do not have many financial resources for extended mission, we cannot afford to continue to do manual tone detection. At the same time, we cannot spend much to automate tone detection. There is additional work to be done to complete the analysis of the experimental tone detection data collected during the primary mission. We will quantify the operational performance of the tone detector and signal characteristics, and possibly improve the detection algorithm with use of a non-linear drift model.

There are scripts on-board the DS1 spacecraft that will send an urgent tone alternating with telemetry during standby (safing) mode. These scripts have been disabled during the primary mission. They will have to be updated, tested, and

uplinked to the spacecraft before Beacon is used for operations.

The Beacon operational concept and the ground visualization software have both been submitted and accepted for NASA Technology Brief Awards. As a result of these reports and other publicity, many external organizations are interested in using and commercializing the Beacon technology.

## 8. CONCLUSIONS

Beacon operations can be viewed as a tool that is valuable in reducing overall mission risk in an environment where decreased tracking is all but mandated by slim operations budgets. It can also be viewed as a technology for conducting low cost mission operations at acceptable risk. The key point here is that NASA policy towards mission risk and cost changed when the visions for smaller, faster, better, and cheaper missions were born. Beacon operations helps enable many more missions with existing tracking resources and is a practical method for minimizing mission risk while decreasing the frequency of telemetry tracking and staffing levels to save operational cost. The Beacon experiment on DS1 has proven the functionality of the technology. It has also shown that it can be effective in reducing downlink volume and frequency, summarizing spacecraft engineering telemetry, and reducing operations costs. Additional use of Beacon on DS1 should prove that Beacon operations and cost reductions are sustainable in the long-term. Future missions should be able to benefit from this proven technology.

## 9. ACKNOWLEDGEMENT

## 10. REFERENCES

Sherwood, R., et. al: 1999, "Lessons Learned During Implementation and Early Operations of the DS1 Beacon Monitor Experiment," *Third International Symposium on Reducing the Cost of Ground Systems and Spacecraft Operations*, Tainan, Taiwan.

Wyatt, E. J., et. al: 1997, "An Overview of the Beacon Monitor Operations Technology," *International Symposium on Artificial Intelligence, Robotics, and Automation in Space*, Tokyo, Japan.

Wyatt, E. J., et. al: 1998, "Beacon Monitor Operations on the Deep Space One Mission," *Fifth International Symposium on Space Mission Operations and Ground Data Systems*, Tokyo, Japan.

DeCoste, D.: 1997, "Automated Learning and Monitoring of Limit Functions," *International Symposium on Artificial Intelligence, Robotics, and Automation in Space*, Tokyo, Japan.

Sherwood, R., et. al: 1997, "Flight Software Implementation of the Beacon Monitor Experiment On the NASA New Millennium Deep Space 1 (DS1) Mission," *Second International Symposium on Reducing the Cost of Spacecraft Ground Systems and Operations*, Oxfordshire, UK.

Chien, S., et. al: 1997, "Resource Scheduling for a Network of Communications Antennas," *Proceedings of the IEEE Aerospace Conference*, Aspen, CO.

Staehle, R.L., et. al: 1996, "Pluto Express: Advanced Technologies Enable Lower Cost Missions to the Outer Solar System and Beyond," *International Low Cost Planetary Missions Conference*, Laurel, MD.

Wyatt, E.J. and J.B. Carraway: 1995, "Beacon Monitoring Approach to Spacecraft Mission Operations," *First International Symposium on Reducing the Cost of Spacecraft Ground Systems and Operations*, Oxfordshire, UK.

## 11. BIBLIOGRAPHY

***Rob Sherwood*** *is a Senior Member of Technical Staff at the Jet Propulsion Laboratory, California Institute of Technology. He holds a B.S. in Aerospace Engineering from University of Colorado at Boulder, an M.S. in Mechanical Engineering from the University of California at Los Angeles, and an M.B.A. from Loyola-Marymount University. Robert has received 5 NASA Achievement Awards for his work in Spacecraft Mission Operations. Rob was the MVACS payload project manager on Mars Polar Lander. He is currently working on several research projects involving Planning and Scheduling technologies and Operations Cost Reduction.*

***Miles Sue*** *is a Senior Telecommunications Engineer at the Jet Propulsion Laboratory. He obtained his BSEE from University of Hawaii and MSEE from University of Southern California. He has been with the Jet Propulsion Laboratory since 1974, where he has contributed in a number of areas: Advanced Deep Space Network (DSN) System Studies, pre-mission studies for deep-space missions, Personal, Mobile, and Aeronautical Mobile Satellite systems, and Spectrum Engineering. Recently, he has conducted studies and flight experiments in developing a weak-signal detection technique to facilitate a new, low-cost mission operations concept for future deep-space missions.*

***Alan Schlutsmeyer*** *has over 30 years' experience in computing; over 20 years at JPL. He has done work in data compression, real time systems, spacecraft flight software, computer graphics and graphical user interfaces. In*

*addition to his work on Beacon, Alan is currently working on JPL's X2000 program to develop a new generation of modular spacecraft.*

**Jay Wyatt** *is the Technical Group Supervisor for the JPL Mission Software Technology Group. This group performs mission autonomy R&D and systems engineering for the JPL Mission Execution and Automation Section. Jay is the Principal Investigator for the Beacon Monitor Operations Experiment on the Deep Space One Mission and is leading research into methods for adaptive mission operations to lower mission cost and decrease the loading on NASA's Deep Space Network (DSN). As an additional duty, Jay manages the JPL Infrastructure & Automation Work Area, which funds R&D in adaptive mission operations, advanced telecommunications protocols, and antenna automation. His AI interests are in monitoring and diagnosis of complex systems.*